

LA-UR-00-1448

Approved for public release;
distribution is unlimited.

Title: Interpreting, Representing and Integrating Scientific
Knowledge from Interdisciplinary Projects

Author(s): Mary A. Meyer
Ray C. Paton

Submitted to: *Theoria et Historia Scientiarum*

Los Alamos
NATIONAL LABORATORY

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the U.S. Department of Energy under contract W-7405-ENG-36. By acceptance of this article, the publisher recognizes that the U.S. Government retains a nonexclusive, royalty-free license to publish or reproduce the published form of this contribution, or to allow others to do so, for U.S. Government purposes. Los Alamos National Laboratory requests that the publisher identify this article as work performed under the auspices of the U.S. Department of Energy. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Interpreting, Representing and Integrating Scientific Knowledge from Interdisciplinary Projects

(To appear in *Theoria et Historia Scientiarum*: The International Journal of Interdisciplinary Studies)

Mary A. Meyer
Statistical Sciences Group
Los Alamos National Laboratory
P.O. Box 1663, MS F600
Los Alamos, New Mexico, USA 87545

Ray C. Paton
Department of Computer Sciences
Chadwick Building, Peach Street
University of Liverpool
Liverpool, England L69 7ZF

Abstract

We describe a coherent, eclectic approach to interpreting, representing, and integrating knowledge from different scientific disciplines or communities of practice. The approach, called ECLECTIC, draws from a complementary blend of ethnological methods, the hermeneutic analysis of domains, and ecology. Our description focuses on the conceptual bases of this approach, its value, and uses, particularly in handling the methodological considerations in the overlapping phases of interpretation, representation, and integration. We give examples from our use of the approach and describe how it handles difficult methodological issues: (1) knowing what questions to initially ask of members of science communities, (2) identifying their states of knowledge, (3) determining the analyst's role, (4) determining how the knowledge may be self elicited by the members themselves, (5) verifying that the interpretation and representation of the knowledge is meaningful to the members, and (6) integrating differing representations from the communities.

1. Introduction

*Et augebitur scientia ...*¹

Anecdotes abound on the number of advances in knowledge and understanding that occur at the interfaces between disciplines. From our experience, we know that interdisciplinary work can include a range of activities such as the production of a volume of papers, the analysis and interpretation of knowledge from various science communities, the performance of project work at interfaces between the sciences, and the conducting of ethnological (cultural anthropological) research. In addition to the activities listed above, interdisciplinary work involving information technology may be focused on the production of such artifacts as virtual enterprises, knowledge bases, organizational memories, and expert systems. However, all interdisciplinary work involves interpreting, representing, and/or integrating knowledge from diverse fields and often encounters challenges in these same areas. The work and ideas reported in this paper are based on the need to address these challenges.

¹ "And knowledge shall increase ..." Taken from the book of Daniel 12: 4. Used by Francis Bacon in his *Novum Organon* and often quoted with regard to growth of scientific and technological knowledge.

Our approach incorporates ideas and practices from across a range of disciplines and so we call it ECLECTIC. We believe that it offers a unique perspective on and methodological solutions to the challenges of dealing with knowledge from different disciplines or communities of practice. Communities of practice are informal networks of people who share similar interests, experience, and knowledge (Wenger 1999). These communities form as people pursue shared enterprises over time, such as working on interdisciplinary projects to solve technical problems. People have simultaneous memberships in multiple, overlapping communities. For instance, individuals may belong to different communities of practice based on their affiliation, discipline, current work, membership in professional societies, and positions held within the organization.

In this section, we introduce the concept of “domain,” and define and describe the phases of interpretation, representation, and integration. The concept of domain is central to this paper. We take a domain to be that body of knowledge that a person or community of practice uses to address problems within their sphere of expertise. This description is somewhat problematic in the sense that what a domain contains will not only vary over time but will change according to the make-up of the people involved and in relation to the constraints of the organization in which the people work. Some examples of domains in this volume are as follows:

- applications of mathematical modeling techniques to complex biological problems (Maini).
- the interface between neural dynamics and cognitive synthesis (Sporns).
- roles for time in microphysical processes of biological systems (Matsuno).

The concept of domain underlies the activities of interpreting, representing, and integrating knowledge. We view these activities as overlapping, iterative phases in inter-disciplinary work. The first phase, **interpretation**, is performed with the aim of making the message of text, or spoken language—discourse—understandable to a hearer or reader. Interpretation is more than translation; it is about meaning and context. From the viewpoint of the ECLECTIC approach and the products of its application, meaning within a domain is determined collaboratively. That is, meaning is not only determined by the speaker/writer but also by the listener/reader and the context. Interpretation underlies all interdisciplinary work, whether it is implicit or formal, such as in domain analysis, discourse analysis, hermeneutics and ethnology. In the ECLECTIC approach, domain analysis is formal. We define domain analysis as involving the interpretation and representation of a body of problems and the knowledge that is applied in solving them.

In the ECLECTIC approach, domain analysis requires dialogue between the analyst and the expert in the domain or the member of community of practice. Dialogue is considered necessary from both an ethnological and hermeneutic perspective. Dialogue is a means for the ethnologist to ask the expert not only for answers but more basically for what questions to ask. Dialogue also allows the analyst and the expert to collaboratively negotiate the meaning of knowledge.

In the ECLECTIC approach, **representations of knowledge** are primarily the form in which the experts, in dialogue with the analysts or themselves, communicate their knowledge of a domain. For instance, two examples of representations in this volume are as follows:

- practitioner’s representations of core issues regarding the practice of internal medicine (Dioguardi);

- practitioner's (Perl and Meyer) representations of how to go about doing experimental physics.

Representations are based on the communities' conceptualization of the domain (for example, of how the domain is composed of parts and how these parts relate to one another). In other words, these representations could be considered models. Representations come in a variety of textual and diagrammatic forms, and we give examples from projects in the next section. These representations may not only describe what an expert is thinking in terms of problem solving but also its context and purpose. These representations serve as the basis for communication (*lingua franca*) among members of different scientific fields on a project, particularly for virtual enterprises.

In the literature on artificial intelligence, representations can also refer to the codification of knowledge into forms that are computer operational, such as rules, semantic networks, statistical models, and cases-based reasoning. While this more codified representation is not the focus of this article, we will illustrate it when it follows from the first type of representation, for example, when the experts' natural language descriptions of how a technology will perform given a set of conditions is translated into fuzzy rules for an expert system.

In the ECLECTIC approach, **integration** is generally defined as the coherent combining of disparate sources, types, and levels of information for some enterprise. Integration is needed in a host of situations, including when one has knowledge or representations

- of different types (for example, qualitative and quantitative information);
- of explicit or formalized knowledge and tacit knowledge;
- from different sources (simulations, experiments, observations, subjective judgments);
- from varying levels (data and metadata, and also component, subsystem, and system, micro and macro models such as in simulations, and data, information, and knowledge);
- from different points in time; and
- from different experts; or at its most complex,
- from different communities of practice, who may be geographically separate and communicating electronically.

Often, the purpose of integration is decision-making. The reason that disparate information is being collected and synthesized in the first place is to provide some coherent input to a decision. For example, in a project carried out between an automotive company and the Statistical Sciences Group at Los Alamos National Laboratory (Los Alamos or the Laboratory), the impetus was to predict the reliability and associated uncertainty of automotive products, often while they were still in early development (PREDICT 1999). What the automotive company did with the conceptual design of an automotive product, such as a fuel injection system, hinged on its predicted reliability. If the reliability were too low or the uncertainty too high, the concept would be redesigned. Once it met targeted numbers, the concept was taken to the next phase in its development, manufacturing, or testing. The integration occurred among the automotive communities of design, manufacturing, chemical, and software engineering; between expert judgment and test and warranty data, as these became available; information at the parts, component and subsystem level; and through time.

Integration builds on the phases of interpretation and representation but is greater than the sum of these. In our opinion, integration must not only bring together tangible, meaningful representations of the knowledge but the means for using these representations and the communities who would use them. That is, the means for integrating representations must themselves become integrated across the communities' practices. Because knowledge is not static but evolves through practice, it follows that, unless knowledge is used and updated by its communities, it will become irrelevant. (McNamara's paper on the communities of practice involved in underground nuclear testing eloquently illustrates both the renewal cycles of knowledge and their dissolution). Thus, in our approach, integration of knowledge includes ethnological methods for ensuring that the means for integration is "owned" by the individual or communities and becomes part of their problem solving and/or decision making.

2. Background to the ECLECTIC Approach

In the first subsection, we describe the ethnological perspective that guides the ECLECTIC approach, as well as some of the ethnological methods that we have adapted.

2.1 Description of Ethnology

Ethnology, or cultural anthropology, is the study of cultures; culture is defined as shared knowledge, such as exists in people's mental models and their practices. We believe that ethnology is relevant and useful to the study of interdisciplinary work for four reasons: its focus on cultures, on knowledge as models, on knowledge as practice, and its methods for explicating these.

Ethnology has focused on the study of cultures for over four decades, beginning with exotic third-world countries and recently including corporate or scientific cultures. Recently, ethnologists have come to view scientific disciplines as cultures and hence worthy of study. While ethnologists have recently come to this realization, those working in the sciences may have intuited earlier the cultural nature of their disciplines. That is, they may have been aware that their discipline or community of practice differed from others in terms of its domain of knowledge, traditions and customs, and ways of problem solving.

In ethnology, there have been two modern views of culture—culture as models and culture as practices. We find both of these views to have merit and we use them in the ECLECTIC approach; the former more in the representation phase, the latter more in the integration phase and throughout the approach (especially whenever knowledge is tacit).

In the knowledge-as-models view, culture is defined as *shared knowledge*; that is, it is "not only people's customs and artifacts and oral traditions, but what they must know in order to act as they do, make the things they make, and interpret their experience in the distinctive way they do" (Quinn and Holland 1991, p. 5). What people **know** is their *cultural models*—"the presupposed, taken-for-granted models of the world that are widely shared (although not necessarily to the exclusion of other, alternative models) by the members of a society and that play an enormous role in their understanding of that world and their behavior in it" (Quinn and Holland 1991, p. 4). These cultural models take either the proposition-schematic or image-schematic form, each of which enables different kinds of cognitive tasks to be performed. An example of the former

would be “rules of thumb” or propositions about how one goes about doing “good” research in a particular field; an example of the latter could be a diagram for evaluating the performance of a system. Another aspect of these cultural models is the use of metaphor; people use metaphor to map their knowledge from known physical domains to the unknown physical domains or conceptualizations of social and psychological domains.

In the knowledge-as-practice view, knowledge is defined as “doing” or as competence in some valued practice. Within the context of this present volume of papers, “knowledge as doing” includes the design of automotive engines, creation of sculptures, practice of medicine, or performance of physics experiments. This view of knowledge differs from that in the previous approach, in which knowledge is viewed as cognitive and is not assumed to directly translate to behavior. In this knowledge-as-practice approach, knowledge is not simply in people’s heads but in their interactions with their environment as well. Knowledge is viewed as actions taking place in the stream of experience, which is defined to include “the person’s self, the things in the environment and the factors which provide the background against which the person creates meaning for him or herself” (Kwasnick 1990). Meaning is negotiated as part of the practice; that is, the members individually and as parts of their community interpret phenomena using their individual and the shared frameworks and perspectives of their practice. In the knowledge-as-practice approach, the primary unit of study has become communities of practice—informal networks formed by people as they pursue shared enterprises over time. This approach focuses on the social process by which learning occurs. Individuals learn as they engage in and contribute to the practices of their communities; communities learn as they refine their practices and ensure a new generation of members (Wenger 1999).

Ethnological methods provide an excellent means for studying interdisciplinary research, in particular for explicating the tacit knowledge and the interfaces between members of communities. Tacit knowledge and interpersonal interfaces are elusive to most techniques of study because individuals are typically unable to describe these outside the context of their practice.

One of the cornerstones of traditional ethnological methods is an emphasis on portraying the *emic* view, the ethnologist’s interpretation of insiders’ views in the insiders’ own words. The ethnologist’s interpretation is informed by interviews and observations of the insider. During both the interviews and the observations, the ethnologist strives to record the exact words of the insider and a description of their context for analysis. The emic perspective² and its related practices have been adapted to the ECLECTIC approach.

Another ethnological practice that has been adapted to ECLECTIC involves the use of insiders. In field research, ethnologists have traditionally identified one or two insiders who would inform them about the studied society, such as by explaining the customs, describing their cultural models, and the meanings that their own actions and those of others held. Ethnologists have developed the practice of asking these “informing insiders” how to ask their research questions (Briggs 1986); that is, given that ethnologists may not speak the language nor know which questions are culturally appropriate and permissible to ask, they ask the insiders for guidance.

² We note that the emic perspective is similar to the hermeneutic notion of a shared perspective through shared dialogue, as illustrated in Figure 2.

Once ethnologists have identified insiders, they may employ a variety of interviewing and observational techniques, many of which are of interest to the ECLECTIC approach. For instance, one type of interview comes from the knowledge-as-models school of ethnology; it is designed to elicit the structure of the insiders' knowledge and represent it as taxonomies or rough ontologies (Spradling 1972; 1979). In addition to interviews, observational techniques are used in the ECLECTIC approach to gather information on interactions and patterns in behavior, and generally gain understanding of knowledge in practice. One type of observation that we have adapted to ECLECTIC is participant observation. Participant observation is the classic field technique whereby the researcher is not a detached observer but a participant in the activities of those being studied. The emphasis is on gaining emic insight into how insiders view themselves, their own and others' behaviors. We turn now from ethnological methods to two closely coupled ideas in the ECLECTIC approach—hermeneutics and ecology.

2.2 Toward an Ecological Hermeneutics

Hermeneutics can be described as the analysis of language/text to interpret and represent what another person is saying or writing. In a contemporary setting it reflects a critical approach to the interpretation of text. We expand its scope by treating domains of knowledge and/or expertise in terms of the text metaphor (see below and Paton 1997). The reader should also look at the papers by Erdi and Tsuda, and Lund and Paton in this volume.

We believe that it is reasonable and beneficial to apply hermeneutics to acquiring domain knowledge for the following two reasons: Firstly, domain knowledge is represented textually and graphically, and these representations have historically been the subject of hermeneutics. Secondly, hermeneutics provides a theoretical base and techniques for exploring *meaning and context*, concepts key to domain knowledge. Specifically, hermeneutics portrays meaning as a collaboration and as more than what is said linguistically. Thus, hermeneutics allows the researcher to find meaning in context, and to strike a balance between what experts say and what they mean. For these reasons, hermeneutics is particularly appropriate to performing domain analysis and adding domain knowledge (Meyer and Paton 1992).

Hermeneutics is derived from the Greek word *ermhneuw*, which conveyed the meaning of “explain” or “interpret.” Explanation and interpretation remain two central features of an hermeneut's task. Specifically, two aspects of hermeneutic investigations should be distinguished: (1) study of the principles on which a text is to be understood, and (2) interpretation of a text so that its message is understandable to a hearer or reader. Although it would not be appropriate to review the work of hermeneuticists³ in detail here, selected aspects of their work will be considered in relation to the general scope of this paper. In particular, we note certain contributions from Schleiermacher, Dilthey, and Ricoeur.

Hermeneutical studies have developed in a number of related areas such as theology (biblical interpretation), social philosophy (interpretation of human behavior) and existentialism (the

³ Note: We distinguish an hermeneuticist as someone who develops the discipline of hermeneutics from an hermeneut who carries out (say) an interpretative task.

purpose of human existence). The modern origins of the subject are in the biblical hermeneutics of 17th century German theology. According to Thiselton (1995) the idea of the study of hermeneutics was probably first used by Dannhauer in his *Hermeneutic Sacra* of 1654. However, the major pioneer of the modern discipline is Schleiermacher (1768–1834), whose intellectual roots were in Christian pietism, Romanticism, and Kantian transcendental philosophy (Thiselton 1995). Schleiermacher elaborated a theory of understanding that depended on the interaction between two epistemic poles: on the one hand interpersonal/relational and on the other, critical/comparative. The former pole facilitated creative understanding and the latter, critical knowledge (Schleiermacher 1977). A key question that Schleiermacher sought to address was not how or what we may understand but how the process of understanding a text becomes possible. Dilthey's (1833–1911) work was not only applied to text but more generally to a hermeneutical interpretation of human behavior (that is, a social philosophy). He focused on the uniqueness of the "self" and on the fundamental role played by textual coherence in terms of the relationship between parts and whole. A key exponent of an hermeneutical approach during the twentieth century was Ricoeur. His working definition of hermeneutics is the theory of the operations of understanding in their relation to the interpretation of texts (Ricoeur 1981). The Fregean notions of sense and reference are key to Ricoeur's separation between discourse and language, for whereas discourse relates to reality, language relates to itself. Put another way, terms denote—people refer (see Harré 1986). This dichotomy anticipates the ecological dimension of what ECLECTIC is all about as it highlights the distinction between the semantic closure associated with a formal system and the capacity for a natural system to evolve, adapt, and anticipate. From our viewpoint, domains are natural systems and any formal representation will by its nature be incomplete (see below). For Ricoeur the extremes and often contradictions of explanation and interpretation within hermeneutics are brought together through the relations between metaphor and text. Explanation provides the sense or pattern within the discourse (the relations of the parts), and interpretation deals with reference (the relation of the whole). Interpretation finds the metaphor for the text as a whole. The metaphorical step that is made with regard to knowledge acquisition is to treat a domain as text as well as discourse about the domain as text (see Figure 1).

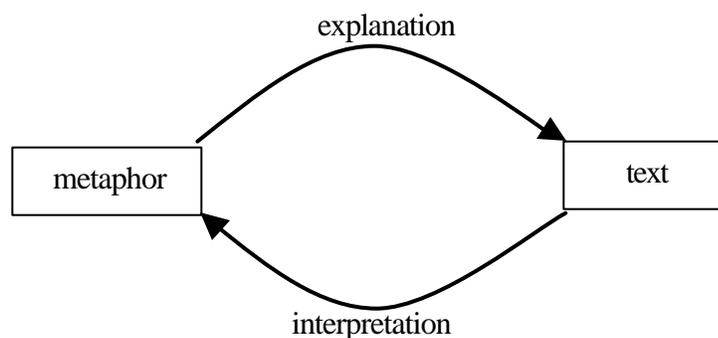


Figure 1. Relation between metaphor and text.

The approach described in this section combines a number of hermeneutical threads. Firstly, there is a hermeneutic emphasis on anticipating and interpreting knowledge. The anticipatory phases in a domain analysis seek to establish the hermeneut's understanding of a domain (that is,

a kind of preunderstanding) and to provide a context in which assumptions and preconceptions are made explicit. This is important for a number of reasons, not least of which is the management of the hermeneut's own thinking and prejudices. Another valuable emphasis of an hermeneutic approach is the focus on dialogue. This can be used as a means of checking a hermeneut's interpretations of knowledge. In order to establish a working rapport between the hermeneut and others, we have developed a number of techniques for probing and characterizing the knowledge—both verbally and visually. Unlike a number of other approaches, these methods are very easy to follow. It can be very difficult to isolate knowledge in a domain and categorize it in any complete or unique way. One solution to this problem is to identify nondisjoint characteristics of a domain. Elsewhere, we have discussed seven such characteristics, namely structure, purpose, theory, metatheory, relation to other domains, history, and metaphor (for example, Meyer and Paton 1992; Paton et al. 1994; Bench-Capon et al. 1996). These characteristics or their anticipation for any domain furnish the hermeneut with a very useful set of analytical and synthetic tools.

As noted in the Introduction, domain was defined as a body of knowledge used by an individual or group of people to address problems. Here we extend the use of the term to include the following domains as they emerge from the dialogical approach:

- cognitive—the understandings of expert and hermeneut,
- dialogue-dependent—the sociolinguistic interactions between expert and hermeneut, and
- referential—the objects to which the expert and hermeneut refer when using words or diagrams or other representations.

Understanding the contexts of the domain is particularly valuable when a range of viewpoints is held among a number of people. The emphasis on interpretation rather than translation is very important and reflects an appreciation of a domain as text. The domain-based approach can provide meaningful dialogue between expert and hermeneut and helps both when it comes to organizing the emerging knowledge in a coherent and usable way. It provides methods for the generation, representation, and containment of a shared perspective recognizing the importance of the evolving meaningful dialogue, as well as the incompleteness of models and languages (summarized in Figure 2). Awareness of constraints on perspectives and models leads to the recognition of a plurality of views that impacts on underlying metaphors as well as ontological issues (Jones and Paton 1997; 1998).

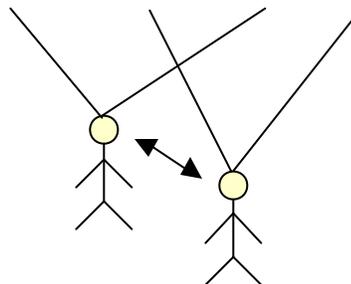


Figure 2. The hermeneutic notion of a shared perspective through a shared dialogue.

Put rather abstractly, we may say that in a dialogue there is exchange interaction between the participating agents whereas in a monological approach there is no interaction per se, rather, one agent acts as an “observer” and defines and describes the “observed” from a singular (idiosyncratic) point of view.⁴ Exchange interaction presumes that observers are also observed although, within our current discussion, the role(s) of expert and hermeneut are dissimilar. It is a powerful way of dealing with complementarity. Elsewhere, and borrowing from certain approaches to microphysics, it is described as an endosystem view. We seek to apply this endosystem view to include context and meaning as products of a shared dialogue. In this case the domain models are generated as a product of dialogue and negotiation, and are not based on the decompositions of an external observer. In some ways this is a reinterpretation of the domain-based view of Meyer and Paton (1992) from the viewpoint of endophysics.

The hermeneutic approach that has been developed for characterizing complex domains of knowledge also has a distinctive ecological flavor (for example, Paton et al. 1993). A domain is not only a text, it is also an ecology, and this latter metaphorical source can be used to account for contexts, interactions, interpretations, history, openness, hierarchy and levels of organization, heterogeneity, and individuality. Domains, like ecosystems, are memory evolutive systems—complex, autonomous systems. (See paper by Ehresmann and Vanbremersh in this volume.) To accommodate this wealth of ideas concerning domains, we need a pluralist approach. A number of metaphorical displacements can be made between domains and ecologies, including systemic and spatial metaphors. As Levins (1984) put it, no single ecological model can simultaneously optimize for realism, generality, and precision. In a similar way, there is no single description of a domain that can be singularly accommodated. This is very evident for example when systemic metaphors associated with ecosystems thinking are examined (for example, Paton 2000). For example, models may access ideas that imply that ecosystems are like

- machines—notably thermodynamic; balance, feedback, input-output, . . .
- circuits—thermodynamic cycles and electrical circuit equivalents (analogues), . . .
- organisms—growth, health, life history, adaptability, evolution, . . .
- economies—role, niche, currency, resource, producer, consumer, . . .
- theatres—stage, play, script, actors, . . .
- societies—interaction, communication, exchange, competition, . . .
- texts—meaning, context, interpretation (for example, plants as phytometers), . . .

What becomes very interesting is that these model types not only access certain vocabularies of words, they also have certain types of diagram associated with them. We also make use of the associations between verbal metaphors and diagrams in developing an appreciation of the ecology of domains.

Domains are open, evolving systems. Within the metaphorical frameworks of domain-as-ecology we may consider how an organism occupies a niche (role/job/function) in an ecosystem, as does an expert within a domain. Different domain experts may fill different niches, and niches can change as the domain changes. Clearly, this is the case in multidisciplinary domains and also within single disciplines. To explore this obvious statement, consider the ecological notion of an *Umwelt*, which was introduced by Uexküll (for example, 1909) to describe those features of the

⁴ Ray Paton acknowledges an important source of this view on monologues and dialogues as coming partly from a conversation with Koichiro Matsuno.

environment that are actually used by an organism. Harré (1990) develops this idea in relation to a realist account of knowledge in scientific domains. An *umwelt* could be described as the species-specific habitat. The expert also has an *umwelt* within the domain in which he/she works. Indeed, the *umwelt* is not only defined by the knowledge and expertise of experts, novices and other workers but also by the organization in which the domain is placed and the history and evolution of this larger “sphere.” This relates to some of the domain characteristics noted above, for example, the relations and common metaphors across domains. In contrast, the *Umgebung* is the prospective habitat of an organism that may change during its life history and over longer time periods. This recognizes that domains have histories that may be linear or cyclical or network-like over time. Correctly anticipating or refuting the broader context of expertise and knowledge need not be and in practice is not cynically viewed within a framework of postmodernist deconstruction. In order to assess particular perspectives, the hermeneutic framework acknowledges a realist account. This places the expert’s tasks and strategies into a context that is cognitive, discursive and organizational.

This realist account is very clearly seen in the role attributed to theoretical knowledge within and between domains. A theory is the evolving cognitive complex that enables people to define objects in the real world as well as model, represent, explain and understand. Metaphors and models are very closely linked. Models based on a source (or sources) that differs from the subject require articulation of one thing in terms of something else. Models that help visualize scientific domains may make use of sources that differ from the subject (for a comprehensive study, see Miller 1996). The language of models is metaphorical. Metaphor provides the linguistic context in which models are described and analogies and similes are made (Harré 1986; Soskice 1985). At the core of a theory is a conception of a mechanism or structure at work. Many theories are more than ordered collections of statements—they also contain an iconic component.

Figure 3 summarizes some interconnections between the major components of the domain of discourse and four other domains. Two of these (organizational and public) are related to the larger organizational concerns and two (personal and modeling) deal with personal/idiosyncratic knowledge. Keeping to the ecological metaphor, the role (niche) played by domain of discourse or dialogue is revealed by the links between domains.

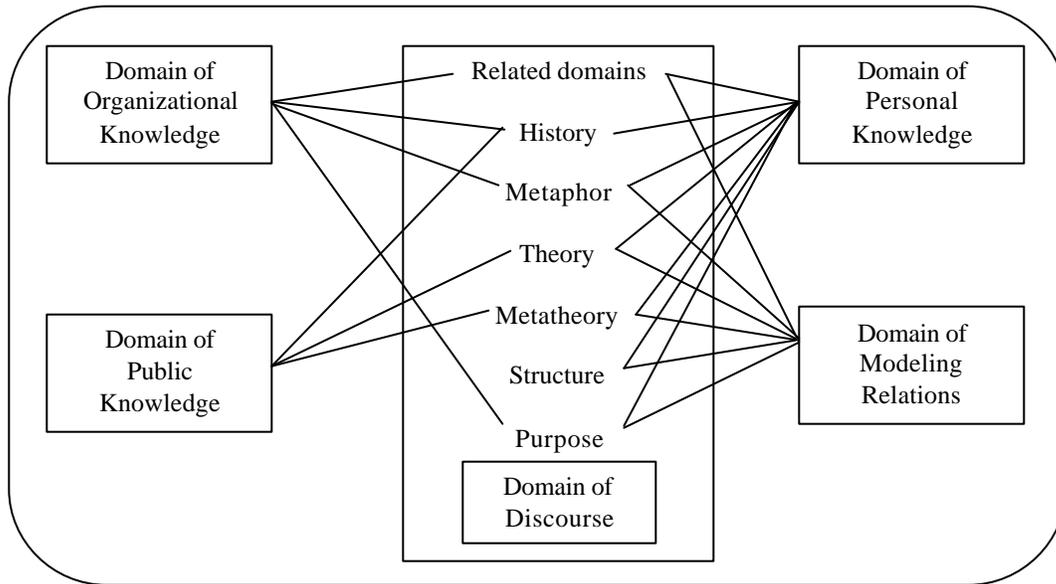


Figure 3. Interrelations with a domain of discourse.

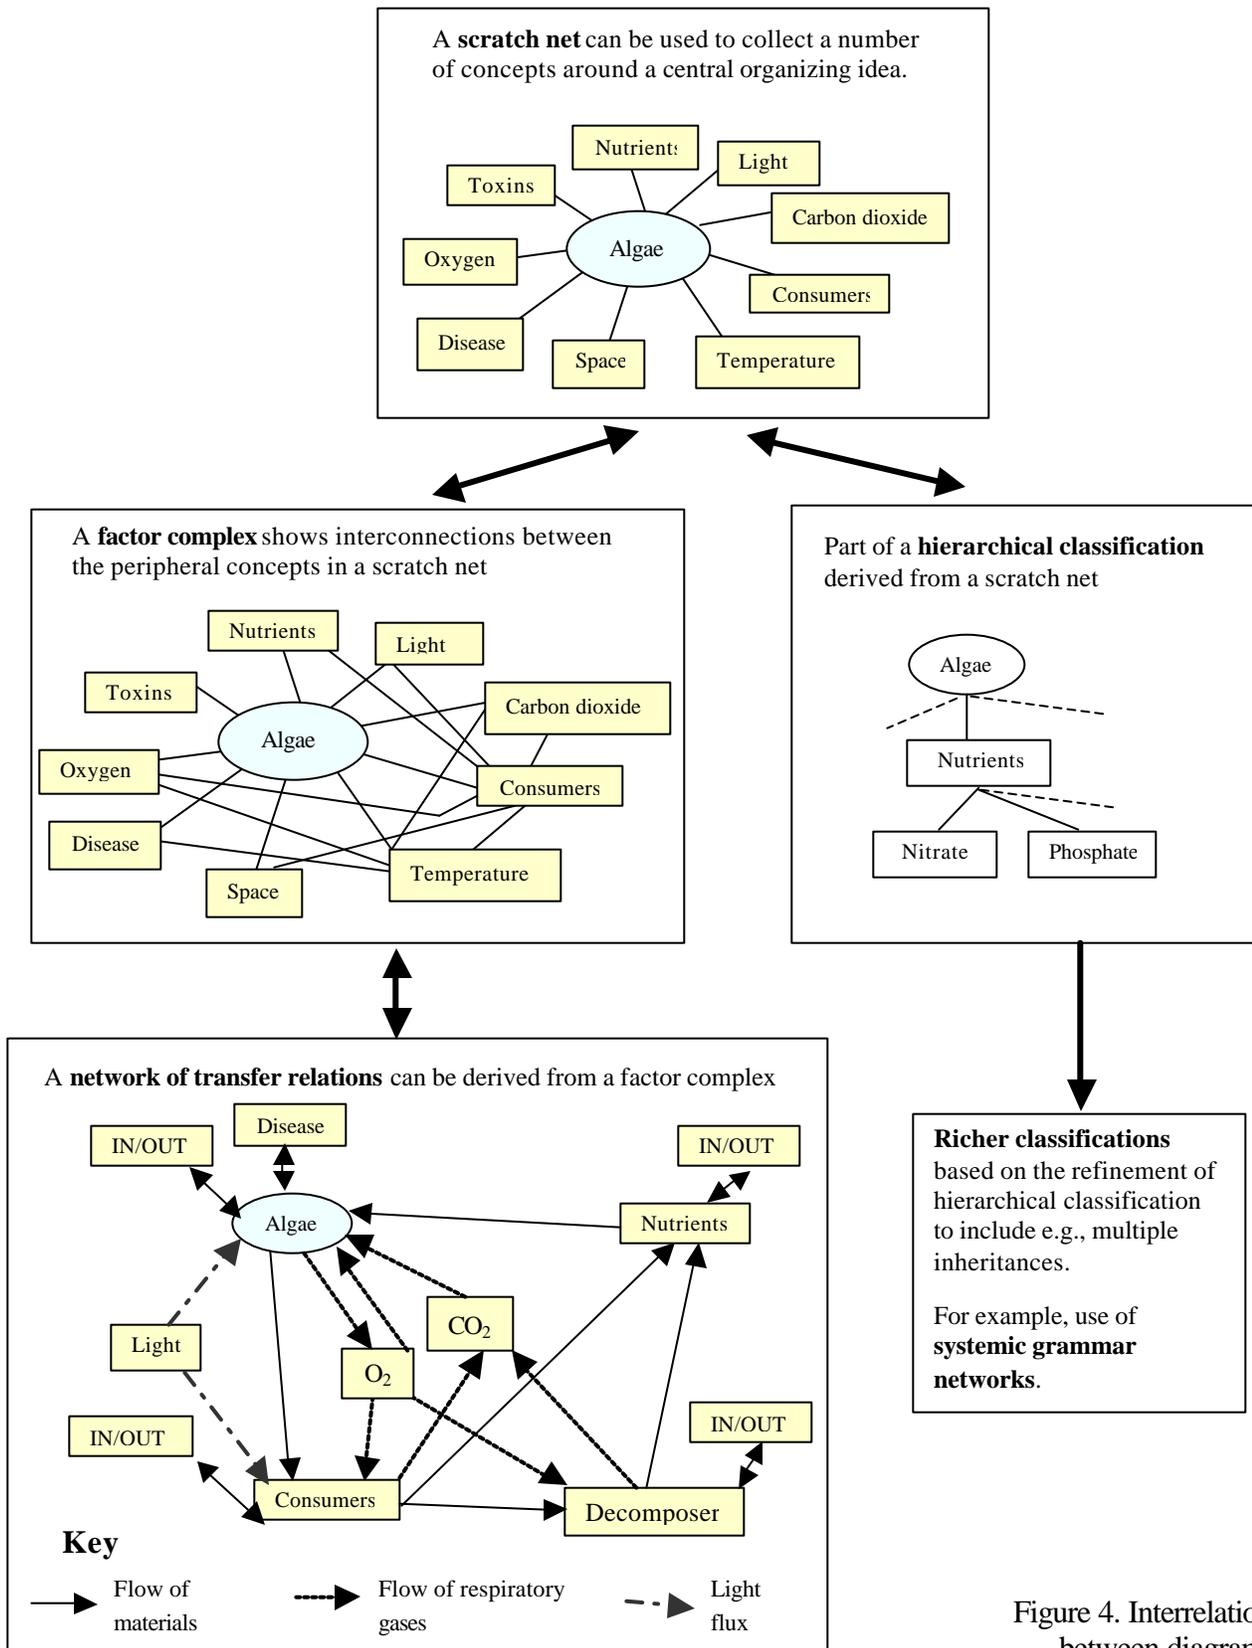


Figure 4. Interrelations between diagrams.

The appreciation of underlying theoretical constructs and the ways in which such constructs are communicated in visual or diagrammatic forms can help an hermeneut explore the structure and functionality of an expert's domain. Figure 4 provides a simple summary of how a number of diagrammatic forms can be interrelated to steer a domain analysis. The scratch net at the top of the figure is a simple method for collecting ideas around a central organizing concept. From this visualization, a number of more detailed diagrams and other forms (such as tables) can be constructed; two examples are shown in the figure. The factor complex reveals interconnections between the peripheral concepts, thus producing a more detailed network. Even more elaborate networks can be derived from this, as shown in the bottom left diagram. Another approach is to refine the scratch net and produce a deeper tree construction from which a classification can be derived. The construction of such hierarchical structures or systemic networks not only provides details about the objects in the domain but also the theoretical constructs used to conceptualize them. Indeed, at this level classifications may be made more complete and underlying taxonomic strategies may be made explicit for the first time (Paton et al. 1994).

3. The ECLECTIC Approach in Practice

Having reviewed some of the general background to the ECLECTIC approach, we now turn to describing the approach in practice. We note that ECLECTIC is an evolving set of tools and techniques, and so we are reporting on an evolving rather than finished piece of (see Paton and Meyer 2000). ECLECTIC is based on the adaptation of the ethnological approach to interdisciplinary projects in Western science. This ethnological foundation is supplemented by the hermeneutic and ecological methods where appropriate. We have organized the next three sections on the phases of interpretation, representation, and integration as follows. We first give a brief description of how we blended the approaches for each phase. We then discuss the methodological considerations in conducting each phase and illustrate these using examples. The examples come from four projects in which various features of the approach have been tested and applied: automotive reliability, radioisotope identification, organizational memory, and weapon reliability.

Automotive Reliability

As mentioned in the Introduction, this project was carried out between a multinational automotive company and the Statistical Sciences Group at Los Alamos. The problems that we were to solve were a combination of statistical and cultural—statistical in that we were to create a process to quantify the reliability, and the related uncertainty, of automotive products while they were still in the concept or design phase, before quantitative test data were available; cultural in that we would be introducing to the automotive company a different way of thinking about reliability, a new way of “doing business.” Our statistical solution involved formally eliciting performance estimates from the product experts, representing these estimates as probability distributions, and integrating these estimates with other relevant information using Bayesian statistics. Our approach to the cultural problem involved ethnological techniques and is described in the next subsection. This project has already met with success in both its statistical and cultural aspects. The process for quantifying reliability has received an R&D 100 award and been adopted for use by the automotive company on all its new concept designs.

Radioisotope Identification

The goal of this project was to create an expert system that would correctly identify radioisotopes from their gamma-ray spectra. Identifying radioisotopes is useful to customs agents or law enforcement officers who must deal with suspicious packages. Our role in this project was to assist the experts in eliciting and representing their own knowledge as rules for the expert system. Experts identify gamma-ray spectra indirectly by the ionization they produce in materials. Measurements of the ionization are recorded as a pulse-height distribution. Because gamma-ray spectra can be measured only indirectly, experts must try to identify imprecise features of the pulse-height distribution and match these to precise features of radioisotope spectra. Our approach in this project was to elicit and observe the process that the experts following in identifying gamma-ray spectra. We then described how the experts could perform a simplified version of the elicitation on themselves to learn their heuristics in identifying spectra.

Organizational Memory

This new project is to capture problem-solving knowledge about projects carried out by the Statistical Sciences Group and to make the knowledge easily accessible and updateable in a web-based repository, a knowledge base. The knowledge base will serve as a means for recording technical lessons learned and as a resource for writing reports and proposals, tracking the progress of projects, and bringing new members up to speed on the projects. The two main challenges are (1) to provide structures for the domain of project knowledge that make sense to a variety of technical staff, such as statisticians, computer scientists, and ethnologists, as well as to managers; and (2) to design the knowledge base and its interface so as to motivate members of the group to contribute to the knowledge. This project is in an early prototype stage in a Lotus Notes Domino application.

Weapon Reliability

In this Los Alamos project, the problem was to quantify the reliability and the related uncertainty of aging nuclear weapon systems. This problem was complicated by the weapon components' potential aging out of their original specifications, but because of the ban on nuclear testing, no new, system-level data were available. The Los Alamos team's approach has been to rely on formally elicited expert judgment when data are sparse or open to interpretation. While this approach is similar to that of the automotive project, there have been greater differences between the weapon communities of practice, in particular between the surveillance engineers and the design physicists. These differences mean greater difficulties in representing and then integrating the communities' reliability assessments. As with the automotive project, performance estimates are represented as probability distributions and integrated using Bayesian statistics. However, the task of integrating the estimates is more complex in the weapons example because of the greater complexity of the system being modeled, the degrees of uncertainty, the differences between the weapon communities, and the time needed to produce data from simulations or experiments.

3.1 Phase 1: Interpretation

In the interpretation phase, ECLECTIC draws on both the ethnological and hermeneutic approaches. The ethnological approach provides guidance in identifying insiders in the community, learning what questions to initially ask them, and determining the sources of explicit

and tacit knowledge. Both the ethnological and the hermeneutic approaches provide means for coming to understand the perspective of the insider through, respectively, interviews or dialogue. (From now on, we will use the term *interviews* when the analyst asks questions and *dialogue* when a more informal style of discourse is taking place.)

We frequently start a project with the interpretation phase—making the knowledge understandable—and our first consideration is identifying *insiders*. As described in section 2.1, insiders are people who will collaborate with us in making the spoken and written communications on the project understandable (to us, to them, and to their communities of practice). We look to insiders as interpreters because the domain being studied and the interdisciplinary communities of practice are often foreign to us, the researchers. The insiders know their domains and their communities of practice. They can provide “entree” into their communities by explaining their workings, identifying and interpreting sources and organizational features of information, and guiding us in eliciting more information and soliciting wider participation. We note that the insider has a much larger role in the ECLECTIC approach than the “friendly expert” did in the early “expert systems” of the 1980s. The role of the friendly expert at that time was restricted to providing rules or cases for the expert system and then checking them.

The first questions we ask the insider are how to ask our questions. That is, when we first enter a project, we may not yet speak the language, know the meaning of certain terms, or even which questions would make sense to the practitioners. For this reason, we explain the aims of the project to the insiders and find out what questions we should ask, and how, and of whom. For example, in using the “asking how to ask” method on the weapon performance project (Meyer and Booker et al. 1999), we learned that the engineers and the physicists think of performance very differently and frame their questions accordingly. For example, the engineers think of performance in terms of measurements, the specifications the system is to meet, and whether it currently meets those specifications, while physicists think in more nebulous terms of processes that could potentially lead to degrees of successful or failed performance.

Typically, we are only certain that we have identified the insiders in retrospect. That is, we describe the role to those who contact us, ask them how to ask questions, and wait to see who begins to carry the work forward in their community (Meyer and Booker et al. 2000). Once the insiders have been identified, we ask them privately what they would personally like to gain from participating in this work and how they will judge its success or failure. We have found that the insiders often have unvoiced reasons, goals, or motivations for championing the work. If we elicit their reasons in advance, we can work toward their realization and ensure the continued participation and good will of the insiders. For example, in the high-tech reliability projects, insiders’ motivations have ranged from wanting to pass on their knowledge before they retired; to developing a process for predicting reliability to the stage where it could be demonstrated in the field, adopted by the company, and applied to all new product development programs (Meyer and Booker et al. 2000). Continued participation is the ideal because it allows the participants to more fully, deeply evolve the interpretation, representation, and integration of the knowledge. In the examples given below, the work on the automotive and weapon reliability projects has been going on for four to seven years.

A second consideration in the interpretation phase is the state of the knowledge. That is, what is the state of the knowledge in the domains? What sources of information are already available at least somewhat codified in text and which remain implicit? To visualize the state of the knowledge, imagine two ends of the continuum. On one end is knowledge represented in the form of text, graphics, or electronic archives. Some examples of this knowledge would be procedures for ways of doing business and accepted scientific theories and practices. This knowledge is fairly static and has been described as explicit or “knowledge as possession.” On the other end is implicit knowledge that has not yet been formalized, recorded in text, nor, often, even articulated. This knowledge resides in the thinking and actions of individuals and their communities. For example, much of implicit knowledge falls into the category of “how to” expertise. This knowledge has been called tacit, expert judgment, or “knowledge as practice.” The first end of the continuum is predominant when the multidisciplinary work is well established, either as a whole or in its individual disciplines. The second end of the continuum dominates when the endeavor is new, with no relevant precursors, such as in designing a novel approach or product.

In our experience, when there is an abundance of formalized knowledge, the role of the analyst will be that of a mirror; that is, the analyst will passively mirror the information and reflect it back. In other words, the analyst adopts the existing terminology. When the state of knowledge is predominantly implicit, the role of the analyst will be to assist in constructing the knowledge, to help bring coherence to it. The analyst’s role will be more active than in the first situation. For example, in the first situation, the analyst will adopt the existing terminology; in the second, the analyst will collaboratively build the vocabulary and negotiate meaning with the other participants.

In the ECLECTIC approach, we determine the state of knowledge by asking the insider: For example, what are the sources of information for this field of inquiry, domain, or project? If there is a lot of information, we ask the insiders to identify what is relevant to the aims of the endeavor and to describe how. For instance, in the weapon reliability project, there was a huge formalized body of knowledge, much of it electronically archived and accessible; but only small portions were declared relevant. This was because the archives documented a former way of quantifying nuclear weapon reliability based on underground testing, a practice that ceased in 1992. The new practices for certifying reliability were not yet fully established at the beginning of this project, and so the majority of knowledge was implicit, despite the magnitude of archived information.

3.2 Phase 2: Representation

Currently, the approach is based mainly on ethnological methods in this phase. However, we expect that hermeneutic representations (Figure 4) will play a greater role in this phase in the future, as we develop techniques that will allow individuals to self elicit and represent their own knowledge. For now, though, we follow ethnological methods in interviewing insiders to learn whether they have existing representations or, if not, in assisting them to create ones compatible with their practices. As in the interpretation phase, the focus in this phase continues to be emic; that is, on portraying the representations from the insiders’ points of view and in their words or diagrams. The ethnological methods are supplemented with techniques from hermeneutics in two situations: (1) use of the container metaphor to assist the insiders in bounding their statements of

the problem; and (2) use of transfer relation diagrams (Figure 4) to depict the flow of information between communities of practice, as this flow has determined by ethnological interviews and observation.

We turn now to describe our current implementation of the approach. Once we have identified the state of knowledge and its sources, we move from a phase that focuses on interpretation to one that includes both interpretation and representation. Representations in the ECLECTIC approach are the form in which the members of the communities and the analysts collaboratively communicate their knowledge of a domain.

If the state of knowledge is fairly well established and written information is available, we work with the insider to make sense of the domain. Our making sense of the domain focuses on defining the scope or boundaries of the domain and its pre-existing structure, particularly the interrelationships of the parts and the purposes, tasks and strategies it addresses. One consideration in the representation phase is the type of representation, and we include and exploit diagrammatic representations as well as purely textual ones. In our experience, practitioners frequently wish to present their knowledge as diagrams (for example, Paton et al. 1994). Often diagrams already exist within their communities of practice or larger organizations that prescribe how tasks are to be performed. Quinn and Holland (1991, p. 5) argue that cultural knowledge may take either proposition-schematic or image-schematic form, each enabling the performance of different kinds of cognitive tasks.

We refer to the project on predicting automotive reliability to illustrate how we represent knowledge in situations in which it has been at least somewhat formalized. We asked our insider to define the domain of product reliability and to identify sources of information on it. In essence, we were asking for the taken-for-granted cultural models (Quinn and Holland 1991) of product engineers that enable them to act as they do and interpret information within the reliability domain. The insider explained that the company approached reliability in terms of meeting performance requirements, for example, for fuel consumption and allowable fuel emissions. Reliability was assessed for each automotive product and the systems, subsystems, and components of which they are composed. For example, Figure 5 (Kerscher et al. 2000) shows a simple generic subsystem D that is composed of components A, B, and C. If the components A, B, and C in the diagram are all in series, the reliability of subsystem D will be the product of the reliabilities of the components.

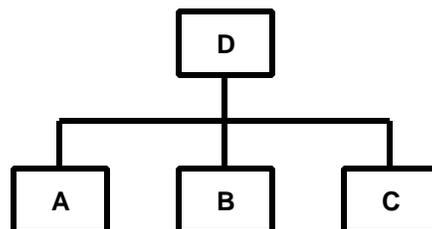


Figure 5. Reliability success tree diagram.

Another consideration in this phase is whether the knowledge must be in quantitative form, which is the norm in scientific applications.⁵ We will describe how tacit knowledge was given quantitative form in the automotive project. In this project, one of the challenges was to quantify the reliability of the product while it was still in the concept or design phase, before quantitative test data were available. We proposed to draw on the knowledge of the product experts and ask them to provide numerical estimates. First, however, we asked the insider how we could ask this question; that is, we asked how the product experts thought about reliability. We learned that the design engineers thought in the metric of incidents per thousand vehicles failing to meet specifications. We also learned that the product experts thought in terms of what could potentially cause the product to fail, or its failure modes. Some examples of failure modes for the outer seal component of a fuel system could be that the seal did not hold or that the extreme temperatures caused the seal to crack.

Using what we had learned, we collaborated with the insiders to draft a set of questions, or a work sheet, to formally elicit initial reliability estimates from the experts. The worksheets named a particular component and asked the experts to identify its significant failure modes, explain why these modes were significant, and estimate their incidence. The worksheets also included three questions to elicit reliability estimates. The first question was essentially “What is your estimate on the most likely number of total failures to occur, for all failure modes, per 1,000 parts?” The second and third questions were similar but asked for the highest number and lowest number of failures the expert reasonably expected. We arrived at this framing of the question through separate dialogues with the automotive insiders and the statisticians on the project team and, later, pilot tests. (In essence, we treated the insiders and the statisticians as two ends on a line and worked the ends to the middle, arriving at a way to frame the question that was satisfactory to both.⁶) The statisticians needed the expert estimates in a form that would translate to uncertainty distributions so they could use a Bayesian approach to the analysis; the insiders and experts needed the questions to match their way of thinking about performance.

The format of most likely number, and the reasonable worst and best case numbers, allowed the statisticians to convert the estimates into uncertainty distributions. Distributions were combined according to the representations of the product as shown in Figure 5. This process will be discussed further in the integration phase.

A practical consideration in this phase involves widening the participation from the insiders to the larger communities of practice. We can illustrate one of these transitions using the automotive project. The measures described below may seem labor and time consuming, and indeed they are, but they have led to the willing participation of the communities—a necessary condition for the project going forward.

After the worksheet was drafted as described above, we asked the insider to pilot test it on himself. We had instructed the insider to think aloud as he completed the worksheet to identify

⁵ Probabilities are often used when experts think in precise numerical form, and fuzzy logic when they do not (Meyer and Booker et al. 2000).

⁶ In the future, we intend to apply the endosystemic view from the ecological approach to such problems. We believe that the endosystemic view might provide insight into how to resolve differences between the communities’ ways of representing their knowledge.

problem areas. The aim of the pilot test was to identify areas where the representation would not fit the communities' reliability practice, ask the insiders for fixes, and then revise the worksheet. (Our procedures are described in detail in Meyer and Booker 1991, pp. 153–159, and are similar to the usability engineering methods of Jacob Nielsen 1993). The insider then selected additional insiders from the design and manufacturing communities to further pilot test the worksheet. We note that the pilot testing took place through teleconferences rather than face to face because the insiders were spread across the continent.

Next, we worked with insiders to ensure that the protocol for administering the worksheets was compatible with the communities of practice. We had learned from the insiders that the product experts worked in teams of four to eight engineers led informally by team leaders. The results of team meetings were generally a consensus on how to proceed in the design or manufacturing of a product, given performance issues. We had also elicited from the insiders the amount of time that the product experts typically spent in meetings assessing product reliability. Taking all this into account, we decided that the protocol for the elicitation would take place in two meetings in which (1) the insider would work with each appropriate team to kick off the effort, request that the team members roughly complete their worksheets individually, and answer any of their questions; and (2) the insider, in cooperation with the team leaders, would use flip charts to elicit and record the team's answers to the worksheet questions. Other parts of the protocol included timing the elicitation to after the company had announced and endorsed the project; and knowing the mechanisms by which the insider initially contacted the teams to solicit their participation. The information the insider gave to the teams to motivate their participation had been carefully designed according to what the insider and what we expected the team members to consider in deciding to participate (for further details, see Meyer and Booker et al. 2000).

Another consideration is the degree to which representation will have to be computer-executable, or automated. While our focus in ECLECTIC has been to evolve representations of the practitioners' thinking, the structured nature of this approach tends to create representations that can be implemented by the computer. In the automotive project, for example, the reliability estimates are propagated (as probability distributions) according to the diagrams of the components, subsystems, and systems to obtain an overall product reliability distribution at a point in time. In the next section, we give additional examples of representations, such as associative networks, that are a first step in representing implicit knowledge and that can be implemented on a computer system.

When the knowledge being studied is largely tacit and embedded in practice, the analyst takes an active role in interpreting and representing "what it is that people know." The analyst may start by asking the insider to identify the domain or to bound the problem. This step may resemble creating a statement of the problem. For example, in creating an expert system to identify radioisotopes, the insider defined the problem and the analyst asked questions until they arrived at a statement of the problem that was acceptable and understandable to both. Much of the dialogue revolved around subtleties of what was being measured directly, namely pulse-height distributions. The final statements of the problem clarified that the instrument was to identify radioisotopes from their gamma-ray spectra but that gamma-ray spectra are detected indirectly by the ionization they produce in materials. Measurements of the ionization are recorded as pulse-height distributions. Thus, experts (and expert systems) must identify the imprecise

features of the pulse-height distribution and match these to precise features of radioisotope spectra (Meyer and Booker et al. 2000).

Another technique for bounding the problem is to introduce the container metaphor from the hermeneutic approach. The analyst can describe this concept to the insider to assist in determining what belongs to the domain and what does not, and therefore need not be considered. For example, the goal of one new project was to create an electronic repository, an organizational memory for statisticians, computer scientists, and ethnologists. The analyst elicited an insider's model of the domain, starting with what the organizational memory should contain. As the analyst continues to carry out the domain analysis, the structure and functioning of the domain can be represented in scratch nets, factor complexes, hierarchical trees, and networks of transfer relations (see Figure 4). We note that the analyst is in an active role of helping evolve the implicit knowledge, as opposed to passively mirroring the knowledge, as is the case when the knowledge is already explicit.

We have found that additional methods are necessary for eliciting and representing task-based knowledge. Task-based knowledge should be elicited or observed as it is practiced. This is because the “know how” of experienced members of a community rapidly becomes unconscious and they cannot recall or describe it out of context. There are several methods for eliciting task-based knowledge, depending on whether the tasks involve browsing electronic records. In the radioisotope project, for example, the spectroscopists referred to electronic libraries of spectra to identify radioisotopes. If the task involves electronic libraries, the method of “machine adaptive learning” may be useful. This method induces users' implicit models from their browsing and represents these as associative networks (Bollen 1998); that is, the sites that the user accesses in sequence are linked together, resembling a factor complex (see Figure 4). For instance, if the spectroscopist accessed the libraries in this order—barium 133, iodine 131, xenon 133, iodine, barium, iodine, and barium—the associative network would show barium and iodine closely linked to each other, and the xenon more distantly linked to the iodine.

For tasks that do not involve navigating through electronic records, we use a combination of three methods to acquire the knowledge—verbal protocol, participant observation, and reflective self-elicitation. The verbal protocol technique from psychology (Ericsson and Simon 1984) involves having the practitioners think aloud as they perform a task so that the researcher can trace and document their thinking. The pilot-test situation described earlier is one example that relies on the think aloud technique, and usability tests are another.

We are also using a customized version of participant observation to acquire task-based knowledge. As mentioned in section 2.1, participant observation is a technique whereby the analyst participates in the activities of those being studied with the aim of gaining insight into how insiders view their own and others' problem-solving tasks (Meyer 1991). We have had to tailor this technique because of logistics. In some cases, our projects involve communities in other states or countries, and travel funds are constrained. In other cases, we have limited access to communities' work areas because of security requirements. Instead of participation, we have used our knowledge of situation, gained by interviews, records research, and domain analysis, to form hypotheses on how insiders view their work, the other communities of practice, and our

joint enterprise. We then check our hypotheses against subsequent relevant information and dialogue with the members.

Self-elicitation is a method that we have been developing as part of ECLECTIC. We define it as occurring when experts reflect upon and elicit and diagram their own knowledge concerning some task, typically as they are performing it. In self-elicitation, the expert plays the role of both himself and the analyst. Self-elicitation is not only a method but also a goal; that is, we aim to enable experts to perform their own domain analysis and propose that this capability aids interdisciplinary work toward becoming self-sustaining. We are developing guidelines and software to assist individuals in scoping, defining, structuring, and representing their own knowledge (Paton and Meyer 2000). Self-elicitation relates to the integration of knowledge within communities of practice, as will be described in the next section.

Experts in the automotive project are making the transition to self-elicitation. In the beginning years, we (the Los Alamos team) elicited from the insiders, and the insiders learned first-hand the technique of “asking how to ask” as we jointly developed the reliability questions. They learned to reflect on how they and the design engineers thought about reliability. They participated in the pilot testing of the questions (the worksheet) and led the eliciting of the estimates from the design engineers. The insiders were then ready and willing to adapt the questions to the next three communities of practice—manufacturing engineers, chemical engineers, and software engineers. For example, for the manufacturing engineers, the insiders changed the question from parts per thousand to parts per million, to reflect the manufacturing engineers’ ways of thinking. They pilot tested the questions and conducted the elicitations of the manufacturing engineers on their own. With the software engineers (engineers who developed the software that ran automotive parts), the insiders needed a little assistance from us because this community of practice thought of performance in a radically different way (for example, hours of operating time).

Frequently, representing the knowledge of communities of practice requires diagramming the flow of information. Understanding and representing the flow of information is necessary to determine how the different domains and niches of experts fit together for accomplishing some common purpose. Moreover, these diagrams can serve as guides in the integration phase.

Our technique has been to diagram the information flow at two levels, top down and bottom up, and to draw on ecological models for guidance. The factor complex or network of transfer relations (Figure 4) are convenient forms for representing the information flow. The top-down view is often provided by insiders who hold managerial positions and who can describe the “big picture” (for example, the functions of each community, the knowledge that they hold, and the type of technical questions that are referred to them). The bottom-up view is obtained from specialists and often reflects their ecological niche. Typically, the specialists explain their expertise, the technical questions they ask and respond to, and their sources of information. We note that this view focuses on the scientist’s daily practice and is more accurate than the generalized picture that emerges from the top-down view.

If the enterprise or the community of practice is newly established, the practitioners may be unable to describe the flow of information, especially at an abstract level. In such situations, we

have asked the insiders for a bottom-up view, for their specific roles and tasks. To illustrate what we might receive from a bottom-up interview, we refer to the project on radioisotopes—the project whose goal was to create an expert system for identifying radioisotopes. We asked an insider to describe the information received by human experts and their decisions and were told the following. The experts, gamma spectroscopists, examine the pulse-height distributions of the observed gamma rays, the detector response functions, and libraries of photo peak energies associated with specific radioisotopes. The experts use the pulse-height distribution to identify the observed gamma-ray peaks. The observed peaks are compared to the detector response function to determine if they are consistent with the detector response or are due to statistics or noise. The experts identify all the features in a pulse-height distribution and refer to the library to categorize a peak. An example of an expert's decision is that it "looks like a Bismuth pulse-height distribution because of the observed peaks, and the extra features look like the Compton edges associated with the observed peaks."

We then asked an insider for the inputs and outputs he expected the expert system to have. The expert system was to have essentially the same information flow as the human expert; that is, the insider expected the expert system to (1) distinguish "peak shape" from that which is "not a peak," (2) compare peak energy to library energy using fuzzy membership, (3) tally the peak matches for all isotopes in the library, and (4) determine the best match to identify the isotope.

A practical consideration in representing knowledge, particularly tacit or task-based, is keeping the representations flexible so that they can evolve. For instance, in the project on weapon reliability, the representations for two of the major communities, surveillance engineers and weapon designers, were in flux for several years, in part because the practice was transitioning from one based on underground nuclear testing to one involving a variety of sciences. The representations for the surveillance engineers and the designers had to be different to reflect their different ways of thinking and assessing performance; however, their representations also needed to interface because information was exchanged between the two communities. Thus, when one community's representation changed, it typically affected the others' representations and all the representations had to be redone, taking much time and effort.

3.3 Phase 3: Integration

Integration is a daunting task—not only does it involve interpretation and representation, complex tasks by themselves, but it must synthesize varieties of these into a coherent whole. In interdisciplinary projects for instance, the representations from the different experts and communities of practice, or disciplines, will be diverse. These representations will reflect different theoretical bases, assumptions, and types (qualitative and quantitative) and sources of information (simulation, experiments, observations, or subjective judgment). Adding to the difficulty of integration is the small amount of available guidance on how to perform it. For example, in the artificial intelligence literature concerning ontologies, the problem of how integration should be performed is described as "more or less unsolved" (Pinto and Gomez-Perez

et al. 1999).⁷ Given the general lack of guidance on how to perform integration, this phase of ECLECTIC is necessarily in more of a developmental stage than those of interpretation or representation.

For the integration phase, ECLECTIC draws mostly on ethnology, especially the knowledge-as-practice approach. This approach focuses on communities' practices, how they work separately and together in solving a problem. Such a focus is necessary to understanding how to integrate their representations and to designing technologies to bring together their diverse representations and records. Also necessary is the use of interviews and dialogues to check that members agree with how the integration has been done.

Given that our approach to performing integration is still evolving, we anticipate greater use of the hermeneutic approach in this phase, particularly in the analysis of domain interrelations (Figure 3). Specifically, we expect that explicating the theory, meta theory, structure, and purpose of a domain will bring rigor to determining how different domains may be integrated. In addition, we also believe that the ecological concepts of niche, of *umwelt* and *umgebung*, may be useful in clarifying domain interrelationships.

We turn now to our method of implementing integration. A key consideration has been the creation of an integrating technology, an electronic repository such as a knowledge base or organizational memory, to bring together the disparate information, organize, and represent it, and generally make it accessible for use in problem solving or decision making. We have developed three of these knowledge bases for the automotive and weapon reliability and organizational memory projects mentioned earlier. Based on hypertext, these knowledge bases have graphical user interfaces and bring together, often for the first time, the information and representations of different communities of practice. The information ranges from

- data (individual test results) to knowledge (experts' interpretation of the data),
- explicit to tacit knowledge (for example, formal reports to informal presentations and transcripts of interviews in which expert judgment was elicited),
- qualitative (natural language statements about the problem) to quantitative, and
- historical and current data to hypothetical future cases.

The representations range from textual descriptions of the experts' problem solving to reliability diagrams and statistics. The representations were designed as described in the previous sections through interviews with members of the communities of practice. In the oldest knowledge base, the knowledge is largely explicit, namely reports or presentations, and has been uploaded into the appropriate categories and cross-linked by the analysts. In the newer knowledge bases for the automotive and weapon reliability projects, both tacit and explicit knowledge are included. The

⁷ The word *integration* has held different meanings in the field of artificial intelligence. For example, in the field of ontological engineering, Pinto and Gomez-Perez et al. (1999) have identified the following three uses of integration:

- integration—to build a new ontology by assembling, specializing, or adapting other ontologies already available;
- merge—to build an ontology by merging several ontologies into a single one that unifies them all; and
- use—to build an application using one or more ontologies.

However, each of these three types of integration is considered difficult and the problem of integration, largely unsolved (Pinto and Gomez-Perez et al. 1999, p. 7–2).

experts are self-eliciting, providing the structure of the knowledge and then uploading their files into the appropriate categories in the structure in the automotive knowledge base.

A key issue in integration is when and at what level representations should be unified.⁸ For example, in ontological engineering, one of the questions is whether integration should be performed during the elicitation, acquisition of knowledge, or during the implementation, when the representations are made executable on the computer. While we agree that integration does and should occur throughout an enterprise's life cycle, we propose that the unifying representation should occur late in the process and at the most encompassing level. Because the different communities of practice have different knowledge and ways of thinking, it would be a grave mistake to try to force one representation during the elicitation stages. In both the weapon and automotive reliability projects, the different communities of practice—the physicists and engineers and the design, manufacturing, chemical, and software engineers—give estimates of performance using the representations and language with which they are familiar. In these two projects, the unifying representation that we selected was Bayesian statistics, an analysis strategy. We had earlier determined that this unifying representation would be compatible with the communities' practices, their diverse representations, and the project's goals. After the experts had provided their estimates in their respective forms, these were translated into the probability form that would be handled by the unifying representation.

For example, in the automotive project, the design engineers give their estimates of defects in parts per thousand, and the manufacturing engineers, in parts per million for the components or subsystems within a system for some unit of time (for example, at 12 months, 10,000 miles, or a million hours of operating time). For the integration to occur, the experts must choose some common time reference, such as 12 months. The analysts, or statisticians, then identify probability distributions for these estimates and combine them to produce the distribution that represents the whole component or subsystem. All the distributions of the individual elements are then combined according to the reliability logic flow diagram, shown in Figure 5, to form the distribution of the entire automotive product. The reliability, including the uncertainty, is then calculated at various points in time using a time-predictive reliability model to calculate the long-term performance of the entire product (Kerscher et al. 2000; PREDICT 1999).

An overview of the process that we followed to arrive at a reliability probability distribution is illustrated in Figure 6.

⁸ We credit Sallie Keller-McNulty for bringing this consideration to our attention.

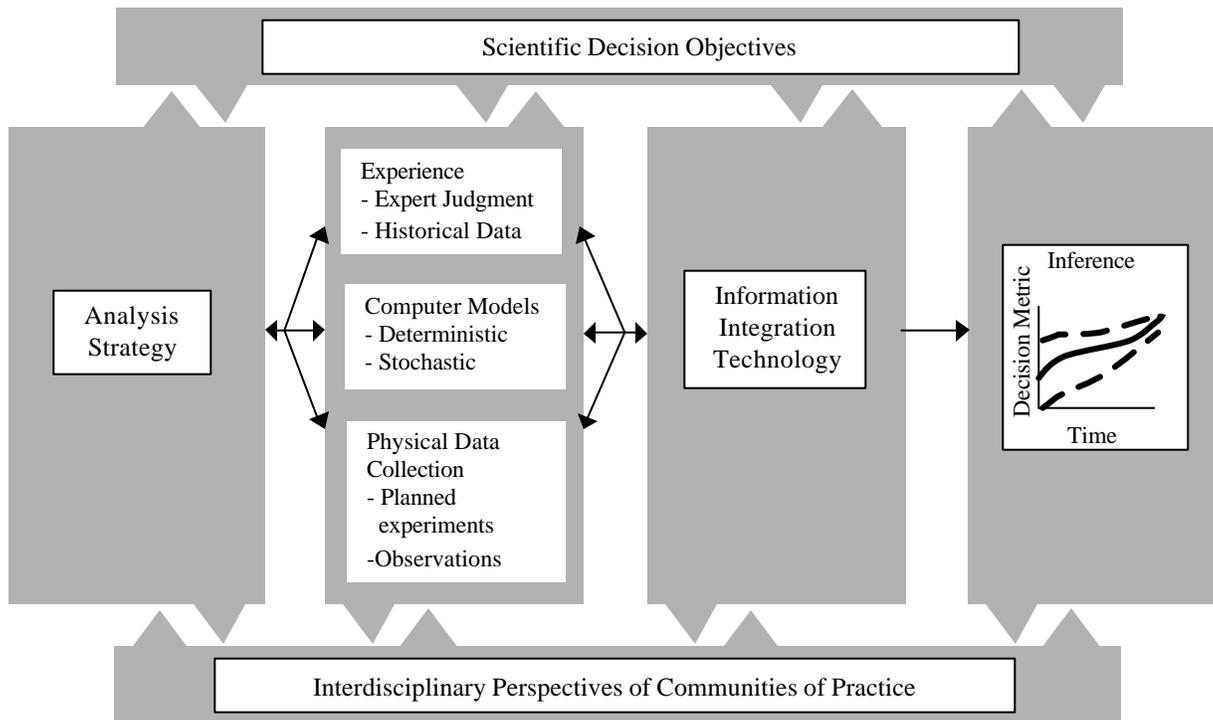


Figure 6. Integration within the communities of practice.

As shown in Figure 6,⁹ we selected an analysis strategy (far left rectangle), in this case Bayesian statistics. As mentioned earlier, this analysis strategy serves as the unifying representation. We selected the analysis strategy according to considerations described earlier in the interpretation and representation phases and depicted in this figure, as the types of information/knowledge (the second rectangle from the left), the interdisciplinary perspectives of the communities of practice and their scientific decision objectives. The representations of the communities' knowledge were brought together via the information integration technology (third rectangle) and unified through Bayesian statistics and a common time reference to provide a decision metric (fourth rectangle). This metric predicts system performance at a point in time in the form of a reliability probability distribution.

As we progress through the integration phase, the dialogue between the analysts and the experts continues to play an important role. Here dialogue serves as a check on the means by which the integration is done, particularly the use of a unifying representation. For example, in the weapon reliability project, the final probabilities and uncertainties for the overall system are reviewed by the communities to ensure that they “make sense.” That is, was there some assumption in the representation that caused the end result to lack credibility, to be overly optimistic or pessimistic? At a more detailed level, the individual experts have reviewed the representations of their judgments and requested changes. For example, some of the weapon experts have requested that labels and caveats be added to the translations of their estimates. Specifically, the plotting of

⁹ We credit Greg Wilson, Sallie Keller-McNulty, and Alyson Wilson for developing this figure.

expert judgments, as probabilities and uncertainties, has caused concern on the part of some experts that decision makers and other users might not be aware of the subjective bases of these judgments. The experts have requested that their plotted results always be labeled as based on *subjective* judgments so that they will not be confused with repeatedly sampled *objective* data.

Another form of dialogue is emerging with the automotive knowledge base; it is between the users of the knowledge base and the knowledge base itself. For example, the experts may use the interface, an application of Lotus Notes Domino software, to electronically self elicit and enter the categories of information (the knowledge structure) that they intend to populate with uploaded files. In this case, the dialogue is between the expert (the disembodied analyst) and the observed knowledge base. Another form of electronic dialogue is occurring among members of the communities of practice. The software allows the different members to view, comment on, or amend each other's entries depending on how the members have defined their own and others' electronic rights. For example, on the automotive project, the core community of practice (the project team) has the greatest voice and electronic access in determining the design and populating of the knowledge base. The project team, in turn, defines the access of the other communities; for example, if they may read only portions of the knowledge base, may edit documents (for example, in completing the worksheets), or may make changes to the structure of the knowledge base. In this way, the communities are able to electronically collaborate on the creation and evolution of the knowledge base. We note that all of the above forms of dialogue can be considered part of the feedback loop, back to when the information is being integrated within the dimension of time.

Time is emerging as an important consideration in integrating knowledge, not only in our projects¹⁰ but also in naturally occurring cases of integration. The classic case of integration is the neurophysiology of the human mind. We mention the neurophysiological example because we have examined it in attempts to extract some of the conditions necessary to all types of integration.

The human mind integrates fragmentary input from all the senses into a coherent, meaningful whole. It is well known that our sense organs gather information and send it to the thalamus, a brain structure, which in turns sends the sensations to regions of the brain's cortex. What is not known is "where and how these fragments of information are integrated into a meaningful whole" (Herbert 2000). Cell physiologists and philosophers of human consciousness have pondered this "binding problem." They have asked how it is that we, for example, can take in the color and boundaries of our offices; the sounds, smells, and textures of the furnishings; and integrate them into a unified sense of "officeness." Llinas, a cell physiologist, has argued that consciousness has more to do with timing than with anatomy and proposed that the thalamus functions as the brain's clock. Using the magnetoencephalograph (MEG) to study the brain's waves, Llinas has shown that "the thalamus is in constant dialogue with the brain's higher processing centers: An electromagnetic loop sends pulses from the thalamus to the cortex, but the different sensory centers of the brain also message the thalamus in return. Consciousness exists when these oscillations are in sync—pulsing at the same rate—so smells, sounds, and so forth assemble in a kind of electromagnetic symphony" (Herbert 2000).

¹⁰ For example, in the automotive project, experts from the different communities had to agree on a common time reference for their representations before these could be unified.

Time seems to be a necessary condition of integration in at least two ways: (1) the synchronization of the different sources of information as mentioned above and (2) the renewal and evolution of the knowledge. As McNamara discussed in her paper, the designing, testing, and refining of prototype nuclear devices at the Nevada Test Site served as a means of integrating the different disciplinary perspectives that the Laboratory's weapon community brought to bear on its tasks: physics, engineering, metallurgy, and the operational expertise necessary to field a test. In preparing for a nuclear test, various communities at the Laboratory worked semi-independently on separate but interrelated problems in fielding a nuclear test. Once preparations had been moved from the Laboratory in Los Alamos to the Nevada Test Site, the test director would rehearse the shot procedures in a series of "dry runs" to check diagnostics, timing and data collection, and firing equipment, ensuring that the individual operations and technologies would function as a unified system when the devices was actually fired. In essence, he was forcing synchronization to a culminating event that would take place in the blink of an eye. This experimental cycle was a proving ground for senior experts and training for the apprentice experts. The code runs, diagnostic results from previous tests, and calculations leading up to and during the test would be the focus of intense discussion between and among the communities. Meanings would be negotiated for every successful or failed step along the way. Individual communities and the weapons community, as a whole, learned through the experimental cycle, so that communal knowledge was constantly evolving.

We note that the Los Alamos knowledge bases differ in the rate at which knowledge is collaboratively updated. We update one Los Alamos knowledge base biannually as a result of feedback and usability tests obtained from users. This knowledge base cannot be updated directly by the dispersed users in a client server mode because they are not connected by a secure network. By contrast, the automotive knowledge base under development is accessed by users directly via e-mail or browser and updated continuously as they refine the structure of the repository and contribute electronic files.

For knowledge to evolve optimally, the technology for integration should facilitate easy, direct, and continuous updating by those most qualified—the communities of practice. We would further argue that if knowledge is to continuously evolve, the technology for integration must be integrated within the community. In other words, the technology, such as the knowledge base, must be "owned" by the communities and become part of their problem-solving practices. (We note that the ECLECTIC approach creates representations that fit the communities' practices, their problem-solving or decision tasks, so that there is one barrier less to their adopting and using the technology.) The knowledge that the community members gain from using the knowledge base to make decisions is folded back into the knowledge base. Similarly, the results of applying any data mining or adaptive learning techniques (for example, to determine implicit mental models of users by the order and frequency with which they access particular sites) also become part of the growing knowledge base.

4. Concluding Remarks

Representing, interpreting, and integrating knowledge present challenging problems to many emerging and established disciplines that must make use of the transfer of ideas, ways of thinking, and practice. This paper, indeed this volume of papers, presents many issues that must

be faced as domains of knowledge evolve. Like the knowledge domains it seeks to address, ECLECTIC is an evolving approach that not only reflects its history (i.e., the background of people applying it and the domains to which it has been applied) but also anticipates and is sensitive to future areas of multidisciplinary knowledge.

References

- Bench-Capon, T. J. M., S. Lynch, and R. C. Paton. (1996). "Principled Development of Knowledge Based Systems—the Importance of Domain History," *Proceedings of the Seventh International Workshop on Database and Expert Systems Applications*, IEEE Press, Los Alamitos, 56–61.
- Briggs, C. L. (1986). *Learning How to Ask: A Sociolinguistic Appraisal of the Role of the Interview in Social Science Research*, Cambridge: Cambridge University Press.
- Bollen, J. (1998). "A Connectionist System to Restructure Hypertext Networks into Valid User Models," *New Review of Hypermedia* 4(1), Sydney.
- Ericsson, K. A., and H. A. Simon. (1984). *Protocol Analysis: Verbal Reports as Data*. Cambridge: MIT Press.
- Harré, R. (1986). *Varieties of Realism: a Rationale for the Natural Sciences*. Oxford: Blackwell.
- Harré, R. (1990). "Exploring the Human Umwelt," in *Harré and His Critics*, edited by R. Bhaskar. Oxford: Blackwell, 297–364.
- Herbert, W. (2000). "Rudolf Llinas: A Grand Unification Theory of the Brain," *U.S. News and World Report*, January 3/January 10, 2000, p. 68.
- Jones, D. M., and R. C. Paton. (1997). "Acquisition of Conceptual Structure in Scientific Theories," in *Knowledge Acquisition, Modeling and Management, Lecture Notes in AI 1319*. Edited by E. Plaza and R. Benjamins. Berlin: Springer, 145–158.
- Jones, D. M., and R. C. Paton. (1998). "Some Problems in the Formal Representation of Hierarchical Knowledge," *Proc. International Conference on Formal Ontology in Information Systems—FOIS'98*, IOS Press series Frontiers in Artificial Intelligence and Applications.
- Kerscher, W. J., J. M. Booker, T. R. Bement, and M. A. Meyer, "Characterizing Reliability During a Product Development Program," European Safety and Reliability (ESREL) 2000, Edinburgh, Scotland, May 15–17, 2000.
- Kwasnik, B. H. (1990). "An Analysis by Means of Naturalistic Analysis of Two Complex Behaviors," in *Proceedings of AAAI Workshop on Complex Systems, Ethnomethodology and Interaction Analysis*, Boston, MA, July 1990, pp. 123–135.
- Levins, S. (1984). "The Strategy of Model Building in Population Biology," in *Conceptual Issues in Evolutionary Biology*, edited by E. Sober, Cambridge: Harvard University Press, pp. 18–27.
- Lund, C. A., and R. C. Paton. (1998). "A Visual Metaphor for Psychoanalytic Training and Supervision," in *Visual Representations and Interpretations*, edited by R. C. Paton and I. Neilson (1999). London: Springer.

- Meyer, M. A., J. M. Booker, R. A. Smith, and T. R. Bement, "Enhanced Reliability Update," *Weapon Insider* 6:1, January/February, 1999, Los Alamos National Laboratory document, Los Alamos, New Mexico, 87545.
- Meyer, M. A., "How to Apply the Anthropological Method of Participant Observation to Knowledge Acquisition for Expert Systems," Los Alamos National Laboratory report LA-UR-91-1876, *IEEE Systems, Man and Cybernetics* 22, no. 5:983–991, September/October 1992.
- Meyer, M. A., and J. M. Booker. (1991). *Eliciting and Analyzing Expert Judgment: A Practical Guide*, Knowledge Acquisition for Knowledge-Based Systems series 5. London: Academic Press.
- Meyer, M. A., and R. C. Paton. (1992). "Towards an Analysis and Classification of Approaches to Knowledge Acquisition from Examination of Textual Metaphor," *Knowledge Acquisition* 4, 347–369.
- Meyer, M. A., J. M. Booker, K. B. Butterfield, W. S. Murray, and R. E. Smith. (2000). "Guidelines for Eliciting Expert Judgment as Probabilities or Fuzzy Logic," Los Alamos National Laboratory report LA-UR-00-218, to be published in *Fuzzy Logic and Probability Applications*, American Statistical Society, SIAM Series.
- Miller, A. I. (1996). *Insights of Genius, Imagery and Creativity in Science and Art*. New York: Copernicus.
- Nielsen, J. (1993). *Usability Engineering*. San Diego: Academic Press.
- Paton, R. C. (1997). "Glue, Verb and Text Metaphors in Biology," *Acta Biotheoretica* 45, 1, 1–15.
- Paton, R. C. (2000). "Systemic Metaphors and Integrative Biology," to be published in *Theoria et Historia Scientiarum*.
- Paton, R. C., M. J. R. Shave, T. J. M. Bench-Capon, and H. S. Nwana. (1993). "Domain Characterisation in Context," in Brezillon, P., *Proceedings of IJCAI Workshop on Context*, France.
- Paton, R. C., S. Lynch, D. Jones, H. S. Nwana, T. J. M. Bench-Capon, and M. J. R. Shave. (1994). "Domain Characterisation for Knowledge Based Systems," *Proceedings of A.I. 94—Fourteenth International Avignon Conference* 1, 41–54.
- Paton, R. C. and Meyer, M. A. (2000). "Toolbox—Method and Enabling Theory for Representing the Structure and Functionality of Knowledge," Los Alamos National Laboratory report LA-UR-00-1002, Los Alamos, New Mexico 87545.
- Pinto, S. H., A. Gomez-Perez, and J. P. Martins. (1999). "Some Issues on Ontology Integration," *Proceedings of the IJCAI-99 Workshop on Ontologies and Problem-Solving Methods (KRR5)* Stockholm, Sweden, August 2, 1999, pp. 7–1 to 7–12.
- "PREDICT—A New Approach to Product Development," 1999 R&D 100 joint entry by Los Alamos National Laboratory and Delphi Automotive Systems. M. A. Meyer, J. M. Booker, T. R. Bement, and W. Kerscher III. Also in *R&D Magazine* 41, p. 161, September 1999.

- Ricoeur, P. (1981). *Hermeneutics and the Human Sciences*. Cambridge: Cambridge University Press.
- Quinn, N., and D. Holland. (1991). "Culture and Cognition" in *Cultural Models in Language and Thought*. Edited by D. Holland and N. Quinn. Cambridge: Cambridge University Press.
- Schleiermacher, F. D. E. (1977). *Hermeneutics. The Handwritten Manuscripts*. Edited by H. Kimmeler. Missoula: Scholars Press.
- Soskice, J.M. (1985). *Metaphor and Religious Language*, Clarendon Press, Oxford.
- Spradling, J. P., ed. (1972). *Culture and Cognition: Rules, Maps, and Plans*. San Francisco: Chandler.
- Spradling, J. P. (1979). *The Ethnographic Interview*. New York: Holt, Rinehart and Winston.
- Thiselton, A. C. (1995). *Interpreting God and the Postmodern Self*. Edinburgh: T & T Clark.
- Uexküll, J. (1909). *Umwelt und Innenwelt der Tiere*, Springer: Berlin.
- Wenger, E. (1999). *Communities of Practice: Learning, Meaning and Identity*. Cambridge: Cambridge University Press.